

OrangeFS

The Next Evolution of PVFS



PVFS2

Brought Us....

- Modular Architecture
- Support for Standard out of the box kernels
- Distributed Meta-data
- Scalability
- Direct MPI Support
- Hardware Independence
- Stateless Operations

OrangeFS ...

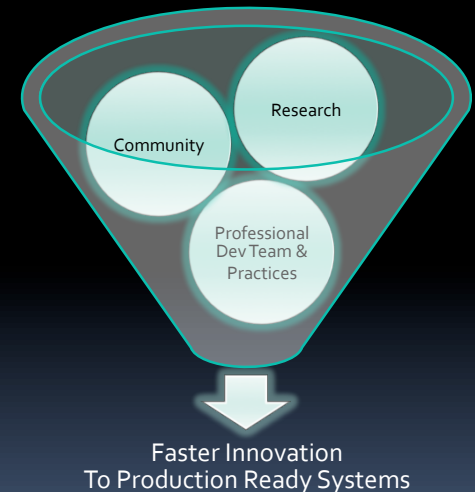
(Current 2.8.4)
(Current 2.8.4)

Added...

- Commercial Grade support (from Omnibond)
- Dedicated Professional Development Team
- Improved stability
- improved meta-data performance for better small file throughput
- replicated data for immutable files
- Supports SSD for Metadata

Still...

- Empower Research
- Community Driven
- 100% Open Source



OrangeFS...


(This Year 2.8.5 & 2.8.6)
(This Year 2.8.5 & 2.8.6)

Will Bring...

- Distributed Directories
 - Based on / Derived from Giga+
- Capability Based Security
- Enhanced Client Libraries
 - Can preload for applications that use standard POSIX calls
 - Client Side Caching (Single System Coherence)
- Windows Client (with Certificate based Authentication)
- WebDav Library (mod_dav_orangefs)
- Enhanced performance Counters and Debugging Tools



Where is OrangeFS Going



Similar Requirements Differing Reasons

	Enterprise / other Research Areas	Exascale
<ul style="list-style-type: none">• Data / IO growth	<ul style="list-style-type: none">• Glimpse of the HEC Space	<ul style="list-style-type: none">• Decade old problem on an exponential path
<ul style="list-style-type: none">• Resilience	<ul style="list-style-type: none">• Up-Time	<ul style="list-style-type: none">• Something is always down
<ul style="list-style-type: none">• Security	<ul style="list-style-type: none">• Shared File Systems	<ul style="list-style-type: none">• Distributed User-base
<ul style="list-style-type: none">• Management	<ul style="list-style-type: none">• Ease of Administration	<ul style="list-style-type: none">• Mass System Management

OrangeFS-Next...


- 128bit UUID Handles
- Data Replication
 - Lazy, On Commit, Continuous
 - Data Tiering
 - Hardware Life-Cycle activities (keep the data online)
- Self Connecting Storage Nodes
 - Ultimately Storage nodes will not need to know about all other nodes (~only ones it shares data with)
 - Nodes will need to understand locality, classes of server, replica info
- Background self healing processes



- ParalleX FS
 - Similar to PGAS (Partitioned Global Address Space)
 - Active Messages (Sending Message Invokes Thread)
 - Futures Synchronization
- Integrate the file system
 - Blur the line between memory and File System
- Highly Object Oriented (C++)
- Exascale

Think of Bucky-Compute-Balls that can be flung to the data it needs

Thanks!



digitalblasphemy.com

Boyd Wilson
b@clemson.edu

